

REMARKS

Reconsideration of the above-identified patent application in view of the amendments above and the remarks following is respectfully requested.

Claims 1-80 are in this case. Claims 41-80 were withdrawn by the Examiner from consideration as drawn to a non-elected invention. Claims 1-8, 10-14, 17, 22, 24-26, 28-31, 34 and 37 have been rejected under 102(e). Claims 9, 15, 16, 18-21, 23, 27, 32, 33, 35, 36 and 38-40 have been rejected under § 103(a). Independent claim 1 has been amended.

The claims before the Examiner are directed toward an interface adapter for a packet network.

In one embodiment of the present invention, the interface adapter includes a memory interface, a plurality of execution engines, a scheduling processor, a plurality of gather engines and switching circuitry. The memory interface is coupled to a memory. The execution engines read, from the memory via the memory interface, work items that have been assigned to them by the scheduling processor, and generate corresponding gather entries. The switching circuitry couples the execution engines to the gather engines so that the gather engines receive the gather entries from the execution engines and generate corresponding packets.

In another embodiment of the present invention, the interface adapter includes a memory interface, a plurality of execution engines, a scheduling processor and a send data unit. The memory interface is coupled to a memory in which are stored work items belonging to transport service instances that are assigned to respective classes of service. The execution engines read the work items from the memory via the memory interface and generate corresponding gather entries. The scheduling processor maintains scheduling queues of the transport service instances according to

the respective classes of service of the transport service instances, and selects transport service instances from the heads of the queues for servicing by the execution engines. The send data unit generates packets responsive to the gather entries.

§ 102(b) Rejections – Tzeng et al. ‘604

The Examiner has rejected claims 1-8, 10-14, 17, 22, 24-26, 28-31, 34 and 37 under § 102(b) as being anticipated by Tzeng et al., US Patent No. 6,912,604 (henceforth, “Tzeng et al. ‘604”). The Examiner’s rejection is respectfully traversed.

Tzeng et al. ‘604 teach a host channel adapter (HCA) **12**, for sending packets on an InfiniBand™ network **10**, in which some link layer operations are performed before the transport layer operations and the rest of the link layer operations are performed after the transport layer operations. Specifically, host channel adapter **12** includes a pre-link module **40**, a transport service module **42** and a post-link module **44**. Pre-link module **40** receives work queue elements (WQEs) from a host and stores the WQEs in a WQE FIFO **50**. A pre-link process module **54** assigns WQEs from WQE FIFO **50** to virtual lane FIFOs **52** according to the WQEs’ service levels. A VL arbitration module **60** sends WQEs from virtual lane FIFOs **52** to transport service module **42** according to the respective priorities of virtual lane FIFOs **52**. Transport service module **42** matches up the WQEs with the corresponding queue pairs (QPs) and sends the WQEs to post-link module **44** that generates the corresponding packets.

Briefly, the difference between the teachings of Tzeng et al. ‘604 and the present invention is that Tzeng et al. ‘604 perform service level classification of WQEs, whereas the present invention performs service level classification of QPs (or, more generally, of transport service instances). HCA **12** of Tzeng et al. ‘604 receives WQEs from the host, classifies the WQEs according to service level, matches up WQEs with QPs and generates corresponding packets. The present invention receives

transport service instances from the host, classifies the transport service instances according to service level (preferably in schedule queues), matches up transport service instances with WQEs (or, more generally, with work items) in execution engines, and generates the corresponding packets in gather engines.

With regard to independent claim 1, the Examiner has identified the “execution engines” recited in this claim with submodules **68** of transport service module **42** of Tzeng et al. ‘604. Claim 1 as now amended describes the execution engines as being

...coupled to the memory interface so as to read from the memory work items corresponding to messages to be sent over the network...
(emphasis added)

In the InfiniBand™ context, this means that the execution engines of the present invention fetch WQEs from the memory, as described in the specification on page 20 lines 26-27:

The assigned execution engine fetches the WQEs of the selected QP from memory 38. (emphasis added)

By contrast, submodules **68** of Tzeng et al. ‘604 receive WQEs from pre-link module **40**. Thus, claim 1 is not anticipated by Tzeng et al. ‘604.

With regard to independent claim 28, this claim, too, recites the limitation that the execution engines are

...coupled to the memory interface so as to read the work items from the memory... (emphasis added)

Furthermore, claim 28 recites the limitation that the scheduling processor is

...adapted to enter the transport service instances to which the work items belong in the scheduling queues according to the classes of service of the instances... (emphasis added)

In other words, in the InfiniBand™ context, the scheduling processor of the present invention classifies QPs according to their service level. As noted above, Tzeng et al.

‘604 classify WQEs, not QPs, according to their service level. Thus, claim 28 is not anticipated by Tzeng et al. ‘604

Furthermore, the present invention, as recited in independent claims 1 and 28 is not obvious from Tzeng et al. ‘604. The overall operational sequence of the present invention includes performing operations, such as classifying QPs according to classes of service, that belong to layers, including the transport layer, that are higher than the link layer, before any link layer operations are performed. Tzeng et al. ‘604 teach specifically against this order of operations, for example in column 3 lines 41-60:

One problem with conventional arrangements for implementing the HCA 12 according to the InfiniBand™ Architecture Specification is that transport layer service would be performed first...Once the transport layer operations have been completed, the packet would be sent to the link layer service...Although this conventional type of implementation has the advantage of precisely following the network layers specified in the InfiniBand™ Architecture specification, such an arrangement requires a substantially large amount of hardware requirements. In particular, the transport layer generally requires more processing power than the link layer because the transport layer involves more complex operations.

and column 8 lines 11-15:

According to the disclosed embodiment, link layer operations are partitioned to enable ordering of work queue entries prior to processing at the transport layer level. Hence, hardware resources, are optimized by prioritizing the work queue entries to be transmitted on the network. (emphasis added)

Thus, independent claims 1 and 28 are allowable in their present form over the cited prior art.

With independent claims 1 and 28 allowable in their present form it follows that claims 2-8, 10-14, 17, 22, 24-26, 29-31, 34 and 37 that depend therefrom also are allowable.

§ 103(a) Rejections – Tzeng et al. ‘604 in view of Pettey et al. ‘712

The Examiner has rejected claims 9, 27 and 32 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of Pettey et al., US Patent No. 6,594,712 (henceforth, “Pettey et al. ‘712”). The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claims 1 and 28 are allowable in their present form. It follows that claims 9, 27 and 32 that depend therefrom also are allowable.

§ 103(a) Rejections – Tzeng et al. ‘604 in view of AAPA

The Examiner has rejected claims 18 and 19 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of Applicant’s Admitted Prior Art. The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claim 1 is allowable in its present form. It follows that claims 18 and 19 that depend therefrom also are allowable.

§ 103(a) Rejections – Tzeng et al. ‘604 in view of Parthasarathy et al. ‘916

The Examiner has rejected claims 15, 20, 21, 23, 35, 36 and 38 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of Parthasarathy et al., US Patent No. 6,831,916 (henceforth, “Parthasarathy et al. ‘916”). The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claims 1 and 28 are allowable in their present form. It follows that claims 15, 20, 23, 35, 36 and 38 that depend therefrom also are allowable.

§ 103(a) Rejections – Tzeng et al. ‘604 in view of Pettey et al. ‘712 and further in view of Parthasarathy et al. ‘916

The Examiner has rejected claim 33 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of Pettey et al. ‘712 and further in view of Parthasarathy et al. ‘916. The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claim 28 is allowable in its present form. It follows that claim 33 that depends therefrom also is allowable.

§ 103(a) Rejections – Tzeng et al. ‘604 in view of Grun ‘591

The Examiner has rejected claim 16 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of Grun, US Patent No. 6,272,591. The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claim 1 is allowable in its present form. It follows that claim 16 that depends therefrom also is allowable.

§ 103(a) Rejections – Tzeng et al. ‘604 in view of Snyder II et al. ‘830

The Examiner has rejected claims 39 and 40 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of Snyder II et al., US Patent No. 6,888,830. The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claim 28 is allowable in its present form. It follows that claims 39 and 40 that depend therefrom also are allowable.

Amendments

An inadvertent typographical error in claim 1 has been corrected. Note that “the host interface” lacks antecedent basis.

In case the Examiner requires support in the specification for this amendment, the Examiner's attention is respectfully directed to Figure 2 that shows execution engines 60 coupled to memory 38 via an execution access arbiter 61 and a translation protection table 62.

In view of the above amendments and remarks it is respectfully submitted that independent claims 1 and 28, and hence dependent claims 2-27 and 29-40 are in condition for allowance. Prompt notice of allowance is respectfully and earnestly solicited.

Respectfully submitted,

Mark M. Friedman
Attorney for Applicant
Registration No. 33,883

Date: October 17, 2006